Survey on the Security of the Quantum ROM

Erica Blum, Makana Castillo-Martin, Michael Rosenberg

December 12, 2019

1 Introduction

The Random Oracle Model (ROM) is a heuristic that has been used to prove the security of hundreds of cryptographic protocols. For over 25 years, it has been used to prove the security of protocols that would otherwise be far more complicated, or even admit no security proof at all. However, since this proof technique is only a heuristic, there is a gap between the notion of security in the ROM and security in in the real world. Over time, the arguments that attempt to bridge this philosophical gap have, in our opinion, become quite strong.

In this paper, we aim to enumerate the arguments in favor of the real-world security of cryptographic schemes proven secure in the ROM, and "port" them to an analogous heuristic, the Quantum Random Oracle Model (QROM), which makes claims about security of schemes against quantum adversaries. The existence of a philosophical argument bridging security in the QROM and security in the real world is similarly important: the QROM has already been used to prove many protocols secure, but it is not self-evident that it is practical as a heuristic.

Our argument for the ROM and QROM will follow in three steps:

1. The (Q)ROM is expressive, and can be used to prove reductions of many varied cryptographic protocols
2. The cryptographic schemes that clearly demonstrate the failure of hash functions to approximate (quantum) random oracles are contrived and unlikely to be constructed in practice.
3. No schemes with security reductions in the (Q)ROM have been shown to have vulnerabilities stemming from the failure of hash functions to approximate a (quantum) random oracle.

We first describe the ROM, the QROM, and their nontrivial relationship.

1.1 Random Oracle Model

The Random Oracle Model was first introduced by Bellare and Rogaway [BR93], and has since been widely used for constructing security proofs of many cryptosystems. Put simply, a protocol in the ROM gives all parties access to a random oracle $\mathcal{O}$. Any party can submit a query $x \in \{0,1\}^*$ to $\mathcal{O}$, and the oracle will return $\mathcal{O}(x) \in \{0,1\}^*$, where each bit is chosen uniformly and independently. If x has been queried before, then $\mathcal{O}(x)$ must be consistent with the previous response. Concrete instantiations of protocols that live in the ROM generally replace the random oracle with a hash function.

The ROM was introduced to address a specific problem in the cryptographic community. At the time, constructions relying on hash functions—for example, full-domain hash (FDH) signatures and the Fiat-Shamir transform—could not be proven secure using any existing techniques. The purpose of the ROM was to provide a heuristic that allowed cryptographers to formally prove the security of these constructions, subject to some assumptions. As Bellare and Rogaway originally argued in [BR93],

> In order to bring to practice some of the benefits of provable security, it makes sense to incorporate into our models objects which capture the properties that practical primitives really seem to possess, and view these objects as basic even if the assumptions about them are, from a theoretical point of view, very strong…We stress that the proof is in the random oracle model and the last step is heuristic in nature. It is a thesis of this paper that significant assurance benefits nonetheless remain.

Indeed, the "last step" refers to the assumption that a hash function ensemble can adequately implement a random oracle, despite the fact that any function with a compact description, by definition, cannot implement a function that returns uniform and independent values for each input. In our argument in favor of the ROM as a useful cryptographic model, we include a hopefully satisfying justification for this admittedly strong assumption.

1.2 Quantum Random Oracle Model

The notion of the Quantum Random Oracle Model first appeared as early as 1996 [BBBV96] but was not formally defined until 2011 [BDF+11]. In security proofs in the QROM, the adversary is given *quantum access* to the random oracle $\mathcal{O}_{\text{quant}}$. In other words, the adversary can send a quantum state $|\psi\rangle = \sum_i \alpha_i |\psi_i\rangle$ to $\mathcal{O}_{\text{quant}}$ and receive the state $\mathcal{O}_{\text{quant}}|\psi\rangle = \sum_i \alpha_i |\mathcal{O}(\psi_i)\rangle$, where $\mathcal{O}$ is a random function. (The honest parties participating in a given protocol are assumed to remain classical.) This fundamentally diverges from the ROM in that a single quantum query can evaluate $\mathcal{O}$ on exponentially many inputs.

This definition of the QROM is motivated by a simple observation: if a protocol uses a hash function as a concrete instantiation of a random oracle, then a quantum adversary could construct a quantum circuit that allows it to query the hash function in quantum superposition. This is fundamentally impossible for a classical computer to do, and thus may afford a quantum adversary a previously unforeseen advantage. Thus, it makes sense to model and generalize this ability with a quantum random oracle.

1.3 Bridging ROM to QROM is Non-Trivial

Quantum-resistant cryptosystems are incredibly important given the imminent reality of quantum supremacy, and many of the best lattice-based constructions for various primitives are shown secure in the ROM. For example, existentially unforgeable signatures and identity-based encryption [GPV08], group signatures [GKV11], and homomorphic signatures [BF11] are all secure and efficient in the ROM. However, other post-quantum cryptosystems that explicitly avoid using a random oracle come at a large cost to efficiency (e.g. signatures [CHKP10,Boy10] and hierarchical identity-based encryption [ABB10,CHKP10]). When designing cryptosystems that are resilient to quantum adversaries, it follows that we should defend against adversaries that can query the random oracle in superposition. Therefore, bridging the ROM and QROM is an important task in this line of work.

If it were the case that security in the ROM implied security in the QROM, then constructing an argument in favor for the practicality of the QROM would be trivial. Unfortunately (or perhaps fortunately), this is not the case.

ROM-Secure $\not\Longrightarrow$ QROM-Secure

Indeed it is not true that a reduction in the ROM necessarily holds in the QROM. [BDF+11] gives a separation result between the ROM and QROM by providing a two-party identification scheme that is secure in the ROM in the presence of an adversary with quantum computation power but insecure in the presence of an adversary that can query $\mathcal{O}$ in superposition.

One may wonder if this degradation of security is the exception or the rule. After all, most proofs of security break down when we change the model from the ROM to the QROM. The following proof techniques are common for showing security in the ROM but are invalidated when an adversary can query the RO in superposition:

Adaptive Programmability A simulator's ability to change the oracle's responses based on previous responses no longer holds when the adversary can query in superposition. The simulator cannot adaptively program the oracle for exponentially many queries at a once.

Extractability A simulator's ability to learn the queries the adversary is interested in and preemptively program the oracle no longer holds when the adversary can hide its true query in superposition.

Efficient Simulation This refers to a simulator's ability to simulate $\mathcal{O}$ efficiently by choosing values on the fly and keeping a hash table of previous queries. But if the adversary queries all possible inputs in

superposition, a poly-time classical simulator cannot efficiently simulate all the (exponentially many) outputs at once.

Rewinding A simulator's ability to reset the adversary to a previous state without detection no longer holds if the adversary has quantum state. If the simulator has performed any measurements to the adversary's state, then un-rewindability follows from the no-cloning theorem.

Fortunately, even though the above key techniques fail to transfer from the ROM to the QROM, many results that are secure in the ROM are still salvageable through other techniques. We describe some of these techniques here:

History-Free Reduction [BDF+11] defines a *history-free* reduction for a cryptographic scheme in the ROM as a reduction to some problem P that essentially does not require introspection into the random oracle queries. The authors formalize this as a framework for the reduction interface. They prove that any reduction that can be expressed this way also holds in the QROM.

Compressed Oracles [Zha18] constructs a method of storing a quantum random oracle as a uniform superposition of truth tables which are "compressed" to just the 0 basis vector in the Fourier basis. Oracle queries thus involve decompressing the database, adding a new query tuple to every database in the superposition, running the query, and then recompressing (which happens to be the same as compressing, since the Fourier change-of-basis operation is an involution). This method of representation allows limited queries on the database itself, such as testing whether $H(x)$ is equal to a particular y. Since not much information is revealed in these queries, they do not perturb the superposition in a noticeable way. The author uses these methods to prove the Merkle-Damgård domain extension construction secure in the QROM.

Lifting via Instability [CMS19] defines and relates different types of games a challenger can play with an adversary: oracle games (where the adversary has access to a random oracle), simulated games (where the adversary has access to a simulation of a random oracle), and database games (where the adversary has access to a simulation of a random oracle, and its precise queries are forgotten). Since the precise queries are forgotten in the database game, it makes sense to think about what queries *could* have been made by the adversary. The authors then define a *database property* as being a predicate on the set of databases, and *stability* as the robustness of a database property under random entry addition. They relate the probability of winning in the classical oracle game to winning the quantum database game and use this technique to prove that Micali's SNARG construction is unconditionally secure in the QROM.

QROM-Secure $\implies$ ROM-Secure

There have been many successful efforts to find sufficient conditions to make this security implication work out. But what about the other direction? Does QROM-security imply ROM-security? The answer, fortunately, is yes. A reduction that is proved in the QROM also holds in the ROM, since a quantum adversary can simply limit their superposition to a single term. We will use this fact in our arguments later on.

We now present an argument in favor of the ROM that we intend to apply just as well to the QROM.

2 The ROM is a Good Heuristic

In this section we attempt argue that the ROM is realistic and useful as a heuristic for security. We include the three points:

1. The ROM is expressive and extensible, having been used to prove the quantum security of many quite different protocols.

2. There are no constructions that are secure in the ROM, but insecure under all concrete instantiations of the quantum random oracle. Or if there is such a construction, it is likely contrived in a similar way as the scheme in [CGH98].

3. No schemes proven secure in the ROM have failed in practice due to a failure in an assumption about the quantum random oracle.

2.1 The ROM is Useful and Expressive

We give examples of varied, important results in cryptography that use the ROM.

Boneh-Lynn-Shacham signatures Signatures are a fundamental primitive in cryptography. Informally, we would like a signature scheme to allow a party Alice to create signatures σ on messages such that (a) any other party can verify that Alice indeed signed that message and (b) no other parties can forge a signature on a message such that it appears to have been signed by Alice.

The Boneh-Lynn-Shacham (BLS) signature scheme was introduced in 2001 [BLS01]. It is known for its simple description and its short signatures.

In the BLS scheme, a prime-order group G and a generator g are given as public parameters. G must be a "Gap Diffie-Hellman" group, which (informally) means that G is a group in which the Decisional Diffie-Hellman (DDH) problem is easy but the Computational Diffie-Hellman (CDH) problem and Discrete Log Problem (DLP) are hard.

A signature scheme is made up of three algorithms: KeyGen, Sign, and Verify. In BLS, Alice generates a secret key and public key by sampling a value x uniformly at random from $\mathbb{Z}_p^*$, where $p = |G|$. Alice's secret key is the value x, and their public key is $\mathsf{pk}_A := g^x$. To sign a message m, Alice computes $H(m)$ and $\sigma = H(m)^x$ (where H is a hash function modeled by a random oracle). Bob can verify a pair (m,σ) by checking that $(g, \mathsf{pk}_A, H(m), \sigma)$ is a valid Diffie-Hellman tuple (i.e., by solving the instance of DDH given by $(g, \mathsf{pk}_A, H(m), \sigma)$).

The Fiat-Shamir Transform Having seen a simple signature scheme, we now move to a slightly more complicated (and much more well-known) technique that has been proven secure in the random oracle model: the Fiat-Shamir transform. Originally presented in the aptly titled paper "How to prove yourself" by Amos Fiat and Adi Shamir [FS87], this technique transforms a 3-step interactive proof-of-knowledge scheme into a non-interactive proof-of-knowledge scheme in the random oracle model. It has been used to create practical implementations of proofs of knowledge and signatures schemes.

Informally, a zero knowledge proof scheme allows a prover to convince a verifier that they know a secret without revealing the secret. Bellare and Rogaway describe the following generic 3-step interactive *zero knowledge* protocol run by a prover P and verifier V [BR93]:

- P has a secret value x and sends an initial message α (sometimes called a commitment) to V.
- V sends a random bit b (sometimes called a challenge) to P.
- P computes a response β based on x and sends it to V, who either accepts or rejects it.

(Additional details not relevant to our discussion have been omitted.)

The Fiat-Shamir transform produces a non-interactive zero-knowledge protocol from this template using a hash function H modeled as a random oracle. Specifically:

- P sends (α, β) to V, where α is a string of P's choice based x and β is computed using α and $H(\alpha)$.
- V either accepts or rejects (α, β).

Note that the challenge bit b is replaced by a call to the random oracle in the form of H.

There is a large body of work studying specific models and security definitions for which Fiat-Shamir can be proven secure, containing both positive and negative results. On the negative side, Goldwasser and Tauman Kalai showed that there is an interactive protocol for which the Fiat-Shamir transform produces insecure signature schemes for any hash function (see below) [GK03]. On the positive side, Bellare and Shoup showed that the transform is secure (for a weaker security definition) in the standard model when the hash function meets certain assumptions [BS07].

2.2 Counterexamples to the ROM are Contrived

In this section we present signature schemes from [MRH03] and [CGH98] which are secure in the ROM, but insecure under any instantiation of the hash function. Such constructions serve as concrete counterexamples to the ROM heuristic. Although in this paper we make the case that the ROM is useful, this is not to say that it is always the best tool for the job! In light of these counterexamples, it is good practice to avoid relying on the ROM when it is not strictly necessary. Nonetheless, the counterexamples we describe below are highly nonstandard, and would never be used in earnest. Some unusual features of these signature schemes include

- Signing algorithms which behave in a deliberately and catastrophically insecure manner for certain inputs
- Message inputs that are parsed as a program tape for a Universal Turing Machine

We give description of the signature scheme of [MRH03], and a more informal description of that of [CGH98], since the latter is quite a bit more complicated in construction and relies on more assumptions than the former to prove security.

The MRH Signature Scheme

The signature scheme $\mathcal{C}$ is built on an EUF-CMA-secure (that is, existentially unforgeable) signature scheme $\bar{\mathcal{C}}$ and an algorithm $\mathcal{D}^{\mathcal{O}}$ which distinguishes whether its provided oracle $\mathcal{O}$ is a hash function f, or a random oracle $\mathcal{R}$. In particular, $\mathcal{D}$ accepts messages of the form (π, t), where π is a program for a Universal Turing Machine, and t is the timeout period for a run of π. $\mathcal{D}$ will test whether $\pi(i) = \mathcal{O}(i)$ for all $i = 1, \ldots, 2|\pi| + k$, where k is the security parameter. If the equality hold for all tested i, $\mathcal{D}$ will return 1. Otherwise it will return 0.

$\mathcal{C}$ is identical in functionality to $\bar{\mathcal{C}}$, except for in its signature algorithm Sign. $\mathsf{Sign}_{\mathsf{sk}}(m)$, will first run $\mathcal{D}^{\mathcal{O}}(m)$. If $\mathcal{D}$ returns 1, the signature algorithm will behave completely insecurely and reveal the secret key sk. Otherwise, the signature algorithm behaves normally and returns the signature under $\bar{\mathcal{C}}$.

The idea of this scheme is that if $\mathcal{O} = f$, an adversary that chooses $\pi = f$ will cause $\mathcal{D}$ to output 1, and thus be able to construct a forgery with high probability. On the other hand, if $\mathcal{O} = \mathcal{R}$, an adversary has negligible probability of being able to produce a π that satisfies the equality for all tested i. Thus, its chances of forging are the same chances it had with $\bar{\mathcal{C}}$, i.e., negligible.

The CGH Signature Scheme

Informally, a relation between inputs x and outputs $\mathcal{O}(x)$ is *evasive* if it is infeasible to find a pair $(x, \mathcal{O}(x))$ that is in the relation (given access to $\mathcal{O}$).

Canetti *et al.* construct their examples using these evasive relations. Specifically, they take an EUF-CMA-secure signature scheme, and then modify it to begin by attempting to find a pair satisfying an evasive relation. If such a pair is found, the modified scheme behaves completely insecurely and reveals sk; otherwise, the modified scheme proceeds as usual.

Since R is evasive, the scheme is secure when $\mathcal{O}$ is a random oracle. When $\mathcal{O}$ is replaced with a element from a function function ensemble $\mathcal{F} = \{f_s\}$ (where s is a seed identifying different functions in the ensemble), and the relation is chosen *based on* $\mathcal{F}$ to be

$$R_{\mathcal{F}} := \cup_k \{(s, f_s(s) | s \in \{0,1\}^k\},$$

then it is easy to find pairs in this relation (because $f_s \in \mathcal{F}$ must be able to be evaluated in polynomial time in order to stand in for the oracle, so an adversary can just compute $f_s(s)$).

It is important to note that this argument only shows that for a given function ensemble $\mathcal{F}$, there exists a signature scheme that is secure in the ROM but not secure when implemented with $\mathcal{F}$. Canetti et al. then use this as a stepping stone to construct a scheme that is secure in the ROM but is insecure under any concrete instantiation of $\mathcal{O}$.

2.3 No Known Real-World Insecure Examples when Secure in ROM

Although it is known that no hash function can truly model a random oracle, there are no non-contrived schemes that "break" in the real world when $\mathcal{O}$ is instantiated with a hash function. In particular, [CGH98] set out to design schemes secure in the ROM but insecure under *any* instantiation of $\mathcal{O}$.

The schemes of [CGH98] violate basic principles of cryptography (e.g. reveal secret information if some rare even occurs). Although [CGH98] concludes that the ROM is not a valid vehicle for proofs of security, they only show this for non-sensible schemes. On the other hand, [KM15] argues that their conclusion should be the exact opposite, that schemes assuming a random oracle do not suggest a real-world weakness. Whichever camp one falls in, we must admit the ROM has provided us with many invaluable schemes for a variety of cryptographic goals (e.g. digital signatures, nonmalleable authentication codes, etc.).

This brings us to the importance of considering a quantum version of the ROM. As quantum computers gain traction and quantum supremacy approaches, it is impertinent to establish quantum-resistant variants of classical cryptosystems, especially the ROM given its prolific presence throughout the crypto world. We will now turn to the quantum analog of the ROM.

3 The QROM is also a Good Heuristic

In this section we attempt to use similar reasoning as the previous section to show that the QROM is realistic and useful as a heuristic for security. We will see that arguing some of these points will be more difficult than in the ROM (given that our confidence in the ROM is based, in part, on a long history of success while constructions in the QROM are still relatively untested) but we proceed with it nonetheless.

3.1 The QROM is Useful and Expressive

As in the first part of the previous section, we argue for the expressivity of the QROM. Despite its limitations, many schemes have been shown to be secure in the QROM.

[BDF+11] proves the QROM-security of a class of ROM-secure signature schemes with history-free reduction. In particular, this technique is applied to the FDH signature schemes of [BR93], [KW03], and [GPV08]. The paper also proves quantum CCA security of a hybrid encryption variant of the public key encryption scheme presented in [BR93].

There is no shortage of literature which use the QROM for their security proofs: [HHK17] uses the QROM to show quantum IND-CCA security of the schemes constructed using the Fujisaki-Okamoto transformation, [Zha18] uses the QROM to show quantum indifferentiability for the Merkle-Damgård domain extender for hash functions, and [Zha12] uses the QROM to prove quantum IND-ID-CPA-security of the identity-based encryption scheme defined in [GPV08].

3.2 Counterexamples to the QROM are Contrived

We argue that the known counterexamples to the QROM heuristic are contrived in the same sense as the known counterexamples to the ROM heuristic. Since it is the case that a reduction in the QROM also holds in the ROM, any counterexample to the QROM would also give a counterexample to the ROM. Since all such constructions in the ROM have been impractical beyond their use as thought experiments, it follows that any such constructions in the QROM are similarly impractical. Since the ROM has existed for over 25 years without practical counterexamples, it inspires confidence in these authors that none exist for it or the QROM.

As of the writing of this document, there is no published construction of a protocol that is proven to be secure in the QROM, and insecure under any choice of hash ensemble. However, the argument from [MRH03] can be easily adapted to the QROM. Since the reduction proof of the signature scheme described in the paper is is history-free, the result from [BDF+11] applies, and we have security in the QROM as well.

Theorem. *If $\bar{\mathcal{C}}$ is a signature scheme which is EUF-CMA secure against quantum adversaries, then $\mathcal{C}$, as defined in [MRH03], is EUF-CMA secure against quantum adversaries in the QROM.*

Proof. We proceed by defining a history-free reduction to the EUF-CMA game on $\bar{\mathcal{C}}$ in the ROM and invoking Theorem 1 of [BDF+11].

Suppose $\mathcal{A}$ is an adversary against the EUF-CMA game on $\mathcal{C}$ with advantage ϵ. We define adversary $\mathcal{B}$ against the EUF-CMA game on $\bar{\mathcal{C}}$ (in the ROM) using the history-free schema. $\mathcal{B}$ is given challenge pk and access to the signing oracle S.

START(pk) : Return pk.

RAND$^{\mathcal{O}_c}(r)$: Return $\mathcal{O}_c(r)$.

SIGN$^{\mathcal{O}_c}(m)$: If $S(m)$ returns sk, goto ABORT(sk). Else, return $\sigma = S(m)$.

FINISH$^{\mathcal{O}_c}(m, \sigma)$: Return (m, σ).

INSTANCE(pk) : Return pk.

Further, we define:

ABORT(sk) : Abort the EUF-CMA game on $\mathcal{C}$ with (m', σ), where m' is a randomly selected message and σ is a signature of m' under key sk.

These interfaces trivially satisfy the properties necessary for a history-free reduction. We now compute the probability that $\mathcal{B}$ wins EUF-CMA on $\mathcal{C}$. The only way for $\mathcal{B}$ to output a forgery is if it got a forgery from $\mathcal{A}$ (which happens with probability ϵ), $S(m)$ returned sk on some m. The likelihood of the latter event happening once is precisely the likelihood that $\mathcal{A}$ (with oracle access to $\mathcal{O}_c$) can efficiently find a π such that $\pi(i) = \mathcal{O}_c(i)$ for all $i = 1, \ldots, 2|\pi| + k$, where k is the security parameter. We can bound this by bounding the probability (over the randomness of $\mathcal{O}_c$) that such a π exists. The set of all programs of length at most ℓ has cardinality less than $2^{\ell+1}$. Thus, the set $\{(\pi(1), \ldots, \pi(q_\ell)) : |\pi| \leq \ell\}$ has cardinality less than $2^{\ell+1}$, where $q_\ell = 2\ell + k$. On the other hand, if we pessimistically assume that $\mathcal{O}_c$ takes only binary values, then the set $\{(\mathcal{O}_c(1), \ldots, \mathcal{O}_c(q_\ell)) : \mathcal{O}_c : \mathbb{N} \to \{0,1\}\}$ has cardinality $2^{q_\ell} = 2^{2\ell+k}$. Then the likelihood that a random oracle has a π of length at most ℓ describing it is

$$\begin{aligned} p_\ell &= \Pr_{\mathcal{O}_c}[\exists \pi : |\pi| \leq \ell, (\mathcal{O}_c(1), \ldots, \mathcal{O}_c(q_\ell)) = (\pi(1), \ldots, \pi(q_\ell))] \\ &= \sum_{|\pi| \leq \ell} \Pr_{\mathcal{O}_c}[(\mathcal{O}_c(1), \ldots, \mathcal{O}_c(q_\ell)) = (\pi(1), \ldots, \pi(q_\ell))] \\ &\leq \frac{2^{\ell+1}}{2^{q_\ell}} = 2^{-\ell-k+1} \end{aligned}$$

Lifting the restriction on the length of π, we see the probability that a random oracle has a program π that describes it on inputs $1, \ldots, q_{|\pi|}$ is

$$p \leq \sum_{\ell=0}^{\infty} 2^{-\ell-k+1} = 2^{-k+2}$$

Thus, by union bound, the likelihood ϵ' of $\mathcal{B}$ winning the EUF-CMA game on $\mathcal{C}$ is at most $\epsilon + Q(k) \cdot p$, where Q is the (polynomial) number of signing queries $\mathcal{A}$ makes. If ϵ' were non-negligible, i.e., $\epsilon' \geq k^{-C}$ for some constant C and all sufficiently large k, then

$$\epsilon \geq k^{-C} - Q(k) \cdot 2^{-k+2}$$

for sufficiently large k, which is non-negligible, and the reduction is complete.

Finally, since we have a history-free reduction which reduces to a problem that is assumed to be hard for quantum-capable adversaries, we have shown that the cryptosystem $\mathcal{C}$ is secure in the QROM.

□

Note that we do not have to assume the existence of quantum-indistinguishable PRFs if we instead replace the PRF in Boneh's proof of Theorem 1 in [BDF+11] with a compressed oracle [Zha18] with a truth table of q many qubits.

The proof showing that the construction of [CGH98] is QROM secure is similarly straightforward, although it does require further assumptions, namely quantum-evasiveness of the constructed relation, and computational soundness of Micali CS-proofs against quantum adversaries.

3.3 No Known Real-World Insecure Examples when Secure in QROM

Similarly to the ROM, there have been no instances of a QROM-based construction failing due to the additional structure imposed by a hash ensemble. To meter this claim, the QROM has only been around for 9 years, and no quantum computers are running any of the algorithms in the literature that have been proven secure in the QROM.

4 Conclusion

The ROM has been used for over 25 years to prove many cryptographic schemes secure. In all of its tenure, it has never been a false friend: no ROM-secure scheme ever ended up being insecure due to poor assumptions. We have hopefully shown in this paper that the QROM should be treated with a similar sense of caution and reliance. It has shown its usefulness time and again in proving security against quantum adversaries, and its security is strictly stronger than that of the ROM. It is our hope that this document has been able to dispel some of the fear, uncertainty, and doubt that may exist around schemes which, in the end, are mathematically false.

References

[ABB10] Shweta Agrawal, Dan Boneh, and Xavier Boyen. Efficient lattice (H)IBE in the standard model. In Henri Gilbert, editor, *Advances in Cryptology – EUROCRYPT 2010*, volume 6110 of *Lecture Notes in Computer Science*, pages 553–572, French Riviera, May 30 – June 3, 2010. Springer, Heidelberg, Germany.

[BBBV96] Charles H. Bennett, Ethan Bernstein, Gilles Brassard, and Umesh V. Vazirani. Strengths and weaknesses of quantum computing. *SIAM J. Comput.*, 26:1510–1523, 1996.

[BDF+11] Dan Boneh, Özgür Dagdelen, Marc Fischlin, Anja Lehmann, Christian Schaffner, and Mark Zhandry. Random oracles in a quantum world. In Dong Hoon Lee and Xiaoyun Wang, editors, *Advances in Cryptology – ASIACRYPT 2011*, volume 7073 of *Lecture Notes in Computer Science*, pages 41–69, Seoul, South Korea, December 4–8, 2011. Springer, Heidelberg, Germany.

[BF11] Dan Boneh and David Mandell Freeman. Homomorphic signatures for polynomial functions. In Kenneth G. Paterson, editor, *Advances in Cryptology – EUROCRYPT 2011*, volume 6632 of *Lecture Notes in Computer Science*, pages 149–168, Tallinn, Estonia, May 15–19, 2011. Springer, Heidelberg, Germany.

[BLS01] Dan Boneh, Ben Lynn, and Hovav Shacham. Short signatures from the Weil pairing. In Colin Boyd, editor, *Advances in Cryptology – ASIACRYPT 2001*, volume 2248 of *Lecture Notes in Computer Science*, pages 514–532, Gold Coast, Australia, December 9–13, 2001. Springer, Heidelberg, Germany.

[Boy10] Xavier Boyen. Lattice mixing and vanishing trapdoors: A framework for fully secure short signatures and more. In Phong Q. Nguyen and David Pointcheval, editors, *PKC 2010: 13th International Conference on Theory and Practice of Public Key Cryptography*, volume 6056 of *Lecture Notes in Computer Science*, pages 499–517, Paris, France, May 26–28, 2010. Springer, Heidelberg, Germany.

[BR93] Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In Dorothy E. Denning, Raymond Pyle, Ravi Ganesan, Ravi S. Sandhu, and Victoria Ashby, editors, *ACM CCS 93: 1st Conference on Computer and Communications Security*, pages 62–73, Fairfax, Virginia, USA, November 3–5, 1993. ACM Press.

[BS07] Mihir Bellare and Sarah Shoup. Two-tier signatures, strongly unforgeable signatures, and Fiat-Shamir without random oracles. In Tatsuaki Okamoto and Xiaoyun Wang, editors, *PKC 2007: 10th International Conference on Theory and Practice of Public Key Cryptography*, volume 4450 of *Lecture Notes in Computer Science*, pages 201–216, Beijing, China, April 16–20, 2007. Springer, Heidelberg, Germany.

[CGH98] Ran Canetti, Oded Goldreich, and Shai Halevi. The random oracle methodology, revisited (preliminary version). In *30th Annual ACM Symposium on Theory of Computing*, pages 209–218, Dallas, TX, USA, May 23–26, 1998. ACM Press.

[CHKP10] David Cash, Dennis Hofheinz, Eike Kiltz, and Chris Peikert. Bonsai trees, or how to delegate a lattice basis. Cryptology ePrint Archive, Report 2010/591, 2010. `http://eprint.iacr.org/2010/591`.

[CMS19] Alessandro Chiesa, Peter Manohar, and Nicholas Spooner. Succinct arguments in the quantum random oracle model. In *TCC 2019: 17th Theory of Cryptography Conference, Part II*, Lecture Notes in Computer Science, pages 1–29. Springer, Heidelberg, Germany, March 2019.

[FS87] Amos Fiat and Adi Shamir. How to prove yourself: Practical solutions to identification and signature problems. In Andrew M. Odlyzko, editor, *Advances in Cryptology – CRYPTO'86*, volume 263 of *Lecture Notes in Computer Science*, pages 186–194, Santa Barbara, CA, USA, August 1987. Springer, Heidelberg, Germany.

[GK03] Shafi Goldwasser and Yael Tauman Kalai. On the (in)security of the fiat-shamir paradigm. In *Proceedings of the 44th Annual IEEE Symposium on Foundations of Computer Science*, FOCS '03, pages 102–, Washington, DC, USA, 2003. IEEE Computer Society.

[GKV11] S. Dov Gordon, Jonathan Katz, and Vinod Vaikuntanathan. A group signature scheme from lattice assumptions. Cryptology ePrint Archive, Report 2011/060, 2011. `http://eprint.iacr.org/2011/060`.

[GPV08] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In Richard E. Ladner and Cynthia Dwork, editors, *40th Annual ACM Symposium on Theory of Computing*, pages 197–206, Victoria, BC, Canada, May 17–20, 2008. ACM Press.

[HHK17] Dennis Hofheinz, Kathrin Hövelmanns, and Eike Kiltz. A modular analysis of the Fujisaki-Okamoto transformation. In Yael Kalai and Leonid Reyzin, editors, *TCC 2017: 15th Theory of Cryptography Conference, Part I*, volume 10677 of *Lecture Notes in Computer Science*, pages 341–371, Baltimore, MD, USA, November 12–15, 2017. Springer, Heidelberg, Germany.

[KM15] Neal Koblitz and Alfred Menezes. The random oracle model: A twenty-year retrospective. Cryptology ePrint Archive, Report 2015/140, 2015. `http://eprint.iacr.org/2015/140`.

[KW03] Jonathan Katz and Nan Wang. Efficiency improvements for signature schemes with tight security reductions. In Sushil Jajodia, Vijayalakshmi Atluri, and Trent Jaeger, editors, *ACM CCS 2003: 10th Conference on Computer and Communications Security*, pages 155–164, Washington, DC, USA, October 27–30, 2003. ACM Press.

[MRH03] Ueli Maurer, Renato Renner, and Clemens Holenstein. Indifferentiability, impossibility results on reductions, and applications to the random oracle methodology. Cryptology ePrint Archive, Report 2003/161, 2003. `http://eprint.iacr.org/2003/161`.

[Zha12] Mark Zhandry. Secure identity-based encryption in the quantum random oracle model. Cryptology ePrint Archive, Report 2012/076, 2012. http://eprint.iacr.org/2012/076.

[Zha18] Mark Zhandry. How to record quantum queries, and applications to quantum indifferentiability. Cryptology ePrint Archive, Report 2018/276, 2018. https://eprint.iacr.org/2018/276.